

FOCUS Release - all supported releases on VM  
June 11, 2001

TM7983

### **Increasing Shared File System Performance**

This technical memo describes techniques that increase Shared File System (SFS) performance under VM FOCUS:

- Define MACHINE XC in the CP DIRECTORY entry for all users (servers and general users). Defining the virtual machine as XC provides better handling of dataspace than defining it as XA or ESA.
- Enable all servers for dataspace. For example:

```
XCONFIG ADDRSPACE MAXNUMBER 50 TOTSIZE 1024M SHARE  
XCONFIG ACCESSLIST ALSIZE 62
```

Defining a FOCUS database in a dataspace can produce significant performance increases. Unfortunately, the server is not forced to use the dataspace.

- Define all directories as DIRECTORY CONTROL rather than FILE CONTROL. DIRECTORY CONTROL directories are faster and can be enabled for dataspace. Refer to your IBM documentation for the implications of this change.
- Update files in dataspace-enabled directories at most once or twice a day.
- Use one of the following VM tools to check on the Working Set Size (WSS) of any SFS virtual machine:
  - SMART
  - RTMESA
  - ESAMON
  - ESAMAP
  - VM:MONITOR
  - VMPRF
  - FCON/ESA
  - CP IND USER *vm\_userid*

RESERVE at least that many pages so that the server does not have to page these in and out all the time.

- Define all servers with OPTION QUICKDSP so that the server is always in queue zero (Q 000) where it will get immediate response from the system.
- Define all servers with OPTION NOMDFS so that the cache controller will give the minidisks on the server more than their fair share of resources.
- Define as many SFS virtual machines as servers as you require for increased performance. A large FOCUS database or a database that has heavy I/O may need to have a dedicated SFS server. This means one SFS server for any database that has critical performance needs. Other SFS servers may serve many non-critical FOCUS databases. The idea is to spread the load across a number of SFS servers.
- Define the SHARE for SFS servers high enough so that the servers can do the work that is required of them. A very large relative share or an absolute value is not unreasonable. Critical SFS servers often specify SHARE RELATIVE 10000 or SHARE ABSOLUTE 5%; the general default value is SHARE RELATIVE 100. If the server does not have enough CPU resources, all users accessing that server will suffer under this constraint. Again, you should use a performance monitor tool to see what is needed. When in doubt, err on the side of greater share.

- Define SFS server minidisks with a large amount of space. Allocate an 800-cylinder minidisk rather than four separate 200-cylinder minidisks. Keep minidisks off volumes with TEMP, PAGE, and SPOOL. SPOOL and PAGE can kill performance.
- SFS servers need a lot of virtual storage, but this will vary based on the load for each server. The following message is a sure sign the server virtual machine size is too small for the load:

DMS5GR3640I File pool server is reclaiming its unused free storage.

Explanation: When the SFS file pool server is running low on virtual storage, it performs a reclaim process to free up additional storage.

System Action: Server processing continues.

Operator Response: None required, but the low storage condition that causes free storage reclaim can be avoided by increasing the server's virtual storage.

System Programmer Response: None.

- When an SFS server starts to run more slowly, it may need to be reorganized. Use the FILESERV REORG command.
- Use VDISK as a TEMP FOCUS work disk:

```
DEFine VFB-512 AS vdev BLK nnnnnnn
```

Issue the HELP CP DEFINE command to see all of the rules. Using VDISK as a TEMP WORK disk will improve performance. Your system may not allow this command, as VDISK increases system paging.

- If a FOCUS database is backed up on a regular basis and has SHADOW set to ON, it might be wise to re-think this setting.

The FOCUS SET SHADOW=ON command causes each record to be written twice, which means twice as many WRITE I/Os as a database without SHADOW ON. Rather than having FOCUS be responsible for the integrity of the database, it would be faster to have CMS be responsible. If an error occurs in the database, it might be better to recover the database from backup instead of keeping duplicate records and generating twice as much I/O. This is a trade-off between performance and absolute database integrity.

The number of READ I/Os is only slightly increased when a database has SHADOW ON.

- You can use the following sample REXX EXEC to COMMIT SFS work:

```
/* EXEC which will COMMIT SFS blocks */
rc = 0
rs = 0
Call CSL 'DMSCOMM rc rs'
if rc /= 0 then say 'DMSCOMM' rc rs
Exit rc
```

Contact Customer Support Service (CSS) at (800) 736-6130 or (212) 736-6130 for a PTF that contains project number 71696. This project includes a DMSCOMM command to commit all work currently open during normal FOCUS operation.

You can issue simple CMS commands in FOCUS by prefixing the command with CMS. For example:

```
CMS CP QUERY TIME
```

The following people are experienced with VM or FOCUS performance have agreed to have their names published in this technical memo:

- Bill Bitner, IBM Performance Guru  
bitnerb@us.ibm.com  
[www.vm.ibm.com/devpages/BITNER/](http://www.vm.ibm.com/devpages/BITNER/)

- Barton Robinson, Velocity Software  
barton@velocitysoftware.com  
www.velocity-software.com
- Dave Kreuter, Vm Resources  
dkreuter@vm-resources.ca  
www.vm-resources.ca